Name_______________ Student ID____________________ Department/Year___________

4th Examination

Introduction to Computer Networks (Hybrid)
Class#: EE 4020, Class-ID: 901E31110
Spring 2021

13:20-14:10 Thursday
Jun 10, 2021

Cautions

1. There are in total 100 points to earn. You have 50 minutes to answer the questions. Skim through all questions and start from the questions you are more confident with.
2. Use only English to answer the questions. Misspelling and grammar errors will be tolerated, but you want to make sure with these errors your answers will still make sense.

1. (ch41, 6pt) Recall the functionalities supported by the network layer. For each of the following statements, tell whether it is true or false.
   (1) Forwarding is one of the supported functionalities. (1pt)
   (2) Breaking of segments into datagrams is one of the supported functionalities. (1pt)
   (3) Forwarding is in the control plane. (1pt)
   (4) Forwardng is in the data plane. (1pt)
   (5) Connection-oriented service is the service model of the Internet network layer. (1pt)
   (6) Connectionless service is the service model of the Internet network layer. (1pt)

Sample Solution:

(1)-(2), (4), (6) are True.

2. (ch42, 6pt) Below is a queuing system at the input port of a router. After table lookup, the incoming packets are classified to 3 different output ports, shown in red, green and blue. There are 12 memory slots in the router. Each memory slot holds a packet. The number on the slot is the packet ID, indicating the order of arrivial.

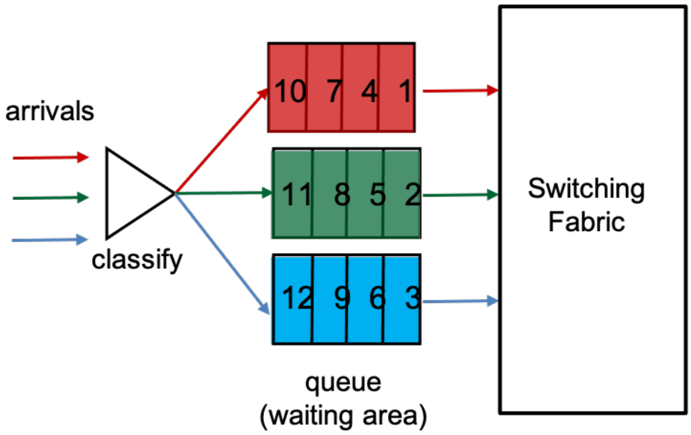

Holding already 12 packets, the system buffer is full. Based on the principle of tail-drop, random (random-drop), and priority (priority-drop) policy. List the packet to drop when a green packet with ID=13 arrives.

(1) Tail-drop policy (1pt)
(2) Random-drop policy (1pt)
(3) Priority + tail-drop policy, with (red, green, blue) queue in (high, medium, low) priority (1pt)

The switching fabric transmits one packet at a time. Based on the principle of priority scheduling, round robin scheduling, and weighted fair queuing, list the packets in the order they will be transmitted further.

(4) Priority scheduling with the (red, green, blue) queue in (high, medium, low) priority. (1%)
(5) Round robin scheduling starting from the red, then green, and finally blue queue. (1%)
(6) Weighted fair queuing with a weight of (red:green:blue) =(3:2:1) and starting from the red queue. (1%)

Sample Solution:

(1) 13
(2) Any of {1, 2, …, 13}

(3) 12

(4) 1, 4, 7, 10, 2, 5, 8, 11, 3, 6, 9, 12

(5) 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12

(6) 1, 4, 7, 2, 5, 3, 10, 8, 11, 6, 9, 12

3. (ch43, 12pt) Recall how fragmentation and reassembly is done when a 4000-byte IP packet runs into a link layer technology that allows MTU = 1500 bytes. The network layer protocol will break the original IP packet into 3 fragments as follows.

(length, ID, flag, offset) = (1500, x, 1, 0)
(1500, x, 1, 185)
(1040, x, 0, 370)

(1) Consider the same 4000-byte IP packet (ID=x) running into a link of MTU = 1980 bytes instead. How many fragments will the network layer send out on the link? (1pt) Tell the (length, ID, flag, offset) of these fragments. (4pt)
(2) Suppose the 4000-byte IP packet run into a link of MTU = 1500 bytes first and then a link of MTU = 900 bytes next. How many fragments will the network layer send out on the MTU = 900 bytes link? (1pt) Tell the (length, ID, flag, offset) of all these fragments. (6pt)

Sample Solution:

(1) 3 fragments

(length, ID, flag, offset) = (1980, x, 1, 0)
(1980, x, 1, 245)
(80, x, 0, 490)

(2) 6 fragments

(length, ID, flag, offset) = (900, x, 1, 0)
(620, x, 1, 110)
(900, x, 1, 185)
(620, x, 1, 295)
(900, x, 1, 370)
(160, x, 0, 480)

4. (ch43, 12pt) Suppose 140.112.42.0/24 is the IP address block assigned to NTU's BL building and 140.112.43.0/24 assigned to NTU's EEII building. The 140.112.42.0/24 block is further divided into 8 groups for 8 office clusters in BL building. The hierarchy is depicted as follows.

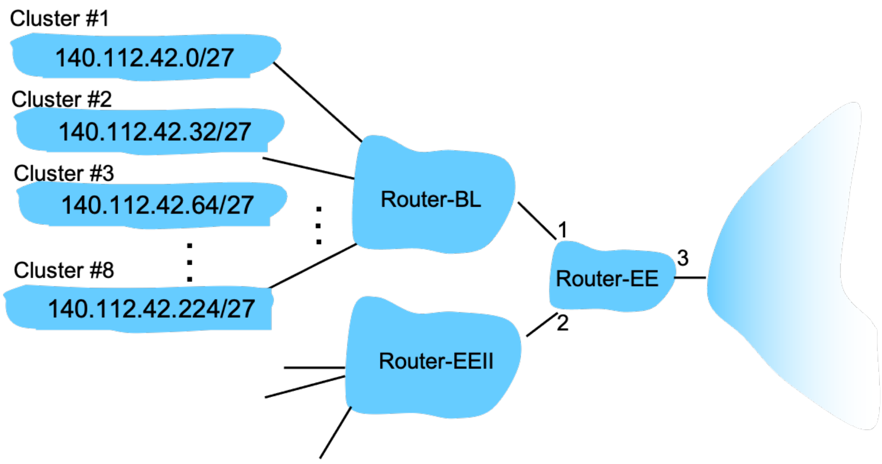

(1) How many subnet addresses will Rouer-BL receive in the route advertisements? (1pt)
(2) How many subnet addresses will Rouer-EE receive in the route advertisements? (1pt)
(3) Write the entries in terms of (prefix, output-port) in Router-EE's lookup table to forward traffic to BL and EEII building. (3pt)
(4) Write the subnet address Router-EE will announce to outside. (1pt)

Suppose cluster #3 originally in BL move to EEII and the staffers ask to keep the old IP addresses to minimize moving effort. The hierarchy is therefore changed as below.

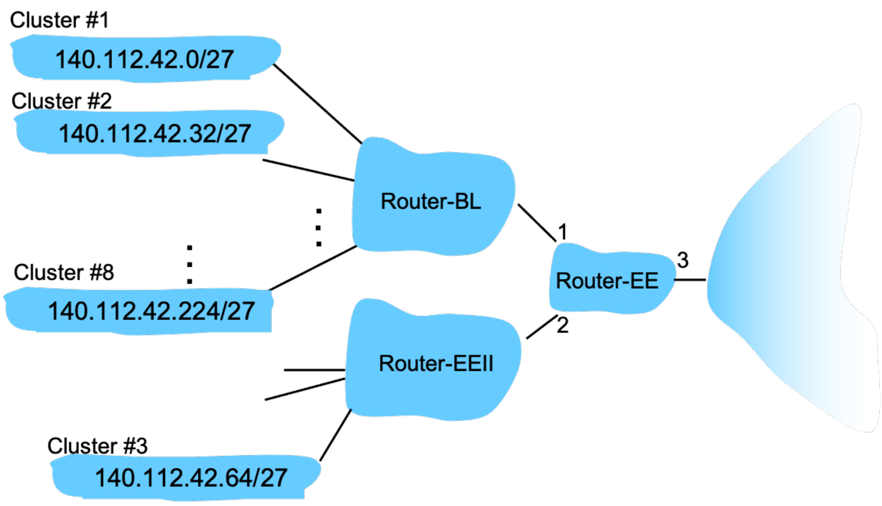

(5) How many subnet addresses will Rouer-BL receive in the route advertisements? (1pt)
(6) How many subnet addresses will Rouer-EE receive in the route advertisements? (1pt)
(7) Write the entries in terms of (prefix, out-port) Router-EE's lookup table to forward traffic to BL and EEII building. (3pt)

(8) Write the subnet address Router-EE will announce to outside. (1pt)

Sample Solution:

(1) 8

(2) 2

(3) (prefix, output-port) =

(10001100 01110000 00101010 ********, 1) for BL and

(10001100 01110000 00101011 ********, 2) for EEII

Or

(140.112.42/24, 1) for BL and

(140.112.43/24, 2) for EEII

(4) 140.112.42.0/23

(5) 7

(6) 3

(7) (prefix, output-port) =

(10001100 01110000 00101010 ********, 1) for BL

(10001100 01110000 00101011 ********, 2) for EEII and

(10001100 01110000 00101010 010*****, 2) for EEII

Or

(140.112.42/24, 1) for BL

(140.112.43/24, 2) for EEII and

(140.112.42.64/27, 2) for EEII

(8) 140.112.42.0/23

5. (ch43, 5pt) Which of the following information does a DHCP server provide?
   (a) IP address
   (b) Network mask
   (c) DNS server IP address
   (d) Lease time
   (e) First hop router IP address

Sample Solution:

(a)-(e)

6. (ch43, 5pt) Which of the following are considered advantages of using NAT.
    (a) A security plus as it is a lot more challenging to scan the local machines
    (b) Supporting the end-to-end argument
    (c) Having access to many local IP addresses when paying for only 1 global address
    (d) No need to reconfigure local addresses when the global address needs to change
    (e) No need to reconfigure the global address when the local addresses need to change

Sample Solution:

(a)(c)(d)(e)

7. (ch43, 4pt) We know some changes can be hard to make, such as making the Internet purely IPv6. We need some serious marketing effort, to motivate various organizations to move on. What would you say or do to get these organizations going?

Sample Solution:

Anything that makes sense

E.g.,

1. compensate the expense for an upgrade to IPv6
2. assign a larger IP address block in exchange of the upgrade
3. say thank you for the service (Mother Teresa's way)
4. hijack all IPv4 routers in the world (Robin hood's way)

(a little hard to pull 4 off and illegal in many places in the world, but it's one way.)

…

8. (ch44, 6pt) The following SDN-based network consists of 3 subnets. Each subnet consists of 2 hosts and 1 router. Let's assume the link costs between each router pair, s1-s2, s2-s3, and s3-s1, are the same. How should the network admin configure the generalzed forwarding tables in s1, s2, and s3 such that each of the 6 hosts gets to send traffic in the shortest distance to every other host?

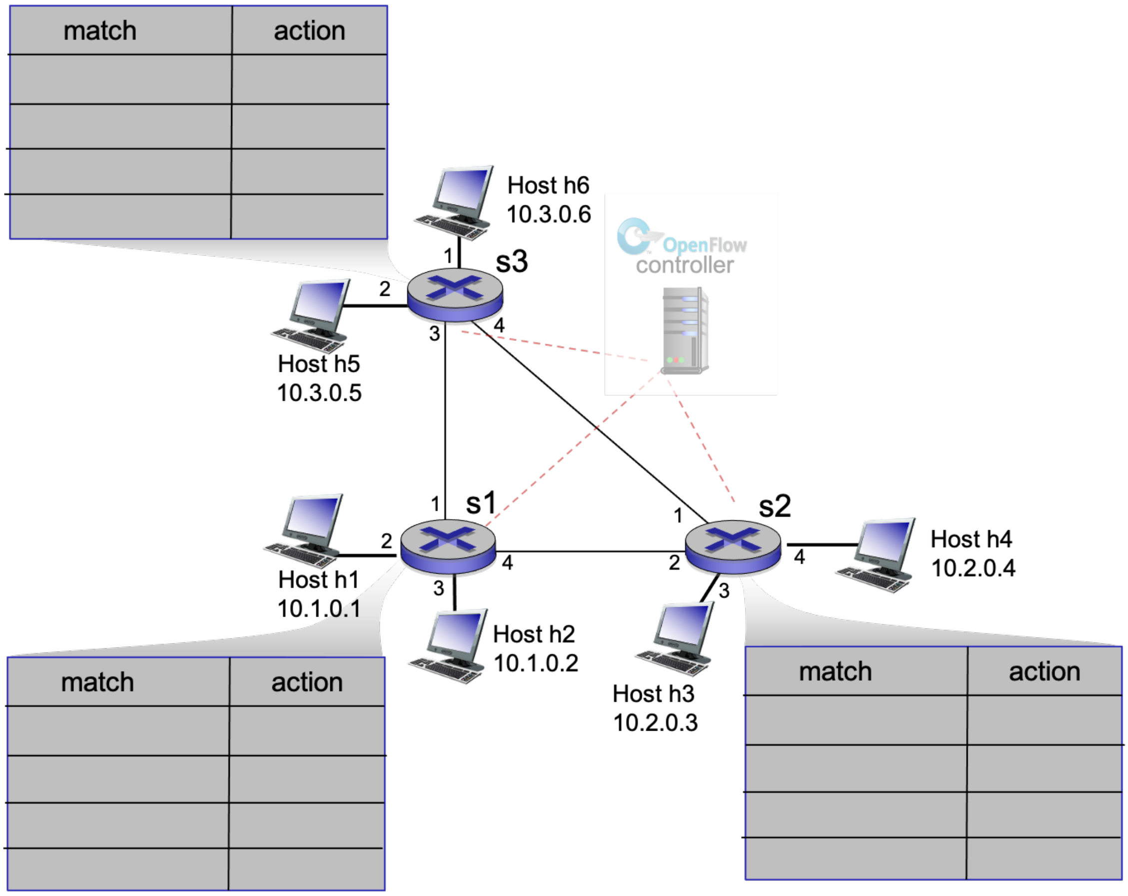

Sample Solution:

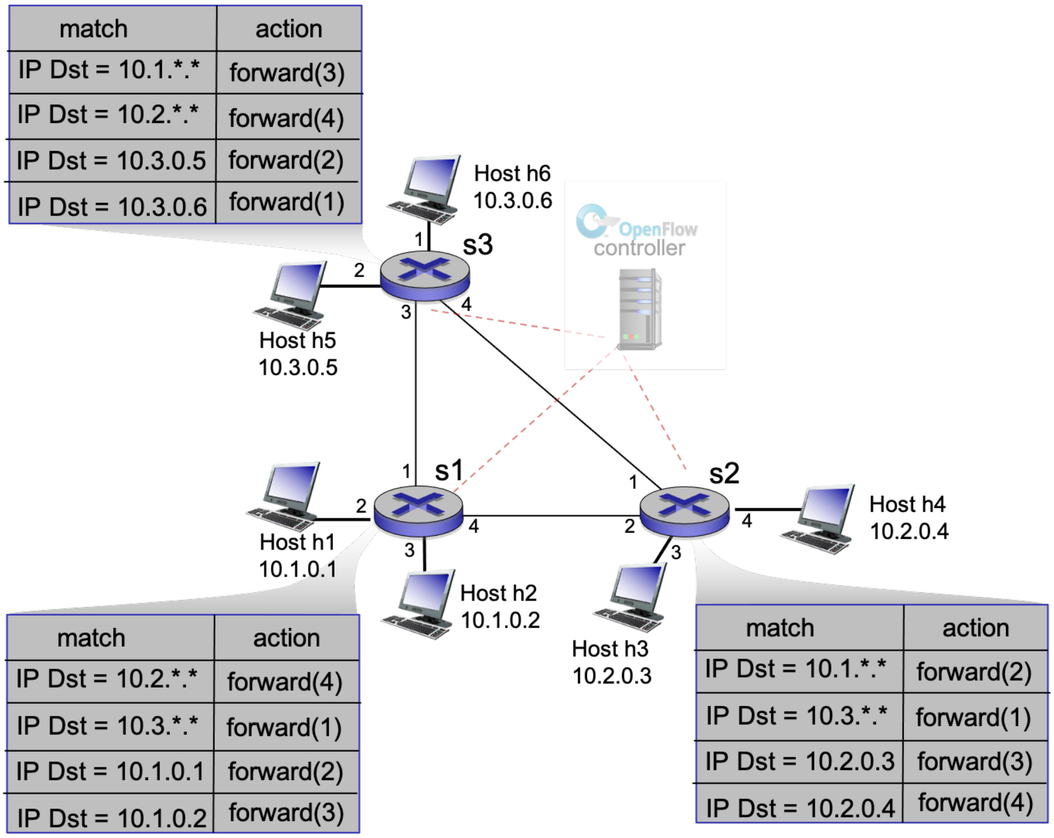

match	action
IP Dst = 10.1.*.*	forward(3)
IP Dst = 10.2.*.*	forward(4)
IP Dst = 10.3.0.5	forward(2)
IP Dst = 10.3.0.6	forward(1)

match	action
IP Dst = 10.2.*.*	forward(4)
IP Dst = 10.3.*.*	forward(1)
IP Dst = 10.1.0.1	forward(2)
IP Dst = 10.1.0.2	forward(3)

match	action
IP Dst = 10.1.*.*	forward(2)
IP Dst = 10.3.*.*	forward(1)
IP Dst = 10.2.0.3	forward(3)
IP Dst = 10.2.0.4	forward(4)

9. (ch52, 10pt) Run Dijkstra algorithm on the following graph. Start from including node u to N’. Iterate through all the nodes and tell the N’ (in order of the nodes being added) upon completion.

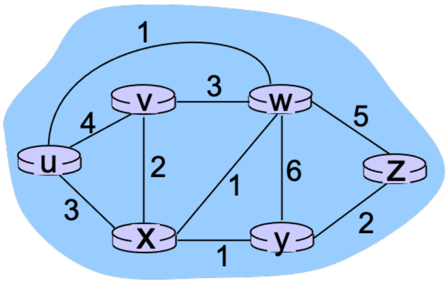

Sample Solution:

N’ = {u, w, x, y, v, z}

10. (ch52, 18pt) We know that good news travel fast in distance vector routing. Let's find out how many iterations it takes to converge when the link cost of (x,y) drops from 3 to 1. To save you some time, Dx, Dy, and Dz have converged to (0, 3, 4), (3, 0, 1), and (4, 1, 0) before the link cost change. This is c0.

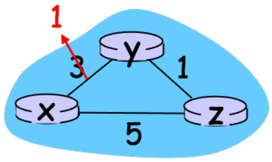

(1) Tell the Dx and Dy after the link cost change. This is c1 (1st iteraton of computation). (2pt)
(2) Upon reacting to the link cost change, node x and y might send distance vector messages to their neighbors. How many distance vector messages will be sent over link (x,y), (y,z), and (z,x) correspondingly? (3pt)
(3) Moving on, suppose the distance vector route computation is optimized such that each node will wait for some time before computation (long enough for the distance vector messages generated within an iteration to arrive). Tell the Dx, Dy, and Dz after receiving the first round of distance vector messages generated from c1. This will be c2. (3pt)
(4) Upon completing the distance vector computation, node x, y, and z might send distance vector messages to their neighbors. How many distance vector messages will be sent over link (x,y), (y,z), and (z,x) correspondingly? (3pt)
(5) After collecting another round of distance vectors from c2, tell the Dx, Dy, and Dz after receiving the second round of distance vector messages. This will be c3. (3pt)
(6) Upon completing the distance vector computation, node x, y, and z might send distance vector messages to their neighbors. How many distance vector messages will be sent over link (x,y), (y,z), and (z,x) correspondingly? (3pt)
(7) How many iterations of route computation take place before the distance vector routes of all nodes converge (such that there is no more distance vector message to send)? (1pt)

Sample Solution:

(1) Right after link cost change (c1), Dx=(0,1,2), Dy=(1,0,1)
(2) 2, 1, 1
(3) After 1st round of DV exchange (c2), Dx=(0,1,2), Dy=(1,0,1), Dz=(2,1,0)
(4) 0, 1, 1
(5) After 2nd round of DV exchange (c3), Dx=(0,1,2), Dy=(1,0,1), Dz=(2,1,0)
(6) 0, 0, 0

(7) 3 iteratons

11. (PA, 16pt) Polly is running 100 servers on port 20001 to 20100 at the PA server. Among all these servers, some are HTTP server (i.e., pollys-PA8-http) and the others HTTPS servers (pollys-PA9-http). The end goal is to scan the ports between 20001 and 20100 and find the port numbers of the HTTPS servers.

    Each of these servers takes an HTTP request (or a TLS encrypted HTTP request) and returns an HTTP response. Check out how telneting and sendng a minimum HTTP request in plain text allows us to tell the two servers apart.

    telnet trial to polly's PA8-http:

```
xdn42o221:PA9 pollyhuang$ ./pollys-PA8-http &
[1] 75953
xdn42o221:PA9 pollyhuang$ Launching server...

xdn42o221:PA9 pollyhuang$ telnet 140.112.42.221 11999
Trying 140.112.42.221...
Connected to xdn42o221.ee.ntu.edu.tw.
Escape character is '^]'.
GET / HTTP/1.1
Host: pollys.ntu.edu.tw

HTTP/1.1 200 OK
Content-Type: text/html; charset=utf-8
Last-Modified: Sun, 06 Jun 2021 12:32:30 GMT
Date: Sun, 06 Jun 2021 12:34:08 GMT
Content-Length: 1087

<pre>
<a href=".DS_Store">.DS_Store</a>
```

    telnet trial to polly's PA9-http:

```
xdn42o221:PA9 pollyhuang$ ./pollys-PA9-http &
[1] 75956
xdn42o221:PA9 pollyhuang$ Launching server...

xdn42o221:PA9 pollyhuang$ telnet 140.112.42.221 11999
Trying 140.112.42.221...
Connected to xdn42o221.ee.ntu.edu.tw.
Escape character is '^]'.
GET / HTTP/1.1
HTTP/1.0 400 Bad Request

Client sent an HTTP request to an HTTPS server.
Connection closed by foreign host.
xdn42o221:PA9 pollyhuang$ 
```

    Exploit how the the HTTP responses are different and program your scanner to identify the HTTPS ports. Now go on to the PA server and log in with the team's username and password. Create a subdirectory `exam4-<student ID>`. Go to the subdirectory and work on the following.

    (1) Create client-e4-1.go such that it is able to send the same minimum HTTP request to

the server running on port 20001 and print the 1st line of the HTTP response from the server. (6pt)

(2) Create client-e4-2.go such that it is able to scan through the servers from port 20001 to 20100 sequentially and print the 1st line of the HTTP response received. (4pt)

(3) Create client-e4-3.go such that it scans through the ports and prints the port numbers of the HTTPS servers. (6pt)

Sample Solution:

Whatever works